

Automatic differentiation of non-holonomic fast marching for computing most threatening trajectories under sensors surveillance

Jean-Marie Mirebeau* Johann Dreot†

April 13, 2017

Abstract

We consider a two player game, where a first player has to install a surveillance system within an admissible region. The second player needs to enter the monitored area, visit a target region, and then leave the area, while minimizing his overall probability of detection. Both players know the target region, and the second player knows the surveillance installation details.

Optimal trajectories for the second player are computed using a recently developed variant of the fast marching algorithm, which takes into account curvature constraints modeling the second player vehicle maneuverability. The surveillance system optimization leverages a reverse-mode semi-automatic differentiation procedure, estimating the gradient of the value function related to the sensor location in time $\mathcal{O}(N \ln N)$.

keywords: Anisotropic fast marching, motion planning, sensor placement, optimization

1 Introduction

This paper presents a proof of concept numerical implementation of a motion planning algorithm related to a two player game. A first player selects, within a admissible class Ξ , an integral cost function on paths, which takes into account their position, orientation, and possibly curvature. The second

*University Paris-Sud, CNRS, University Paris-Saclay

†THALES Research & Technology

player selects a path, within an admissible class Γ , with prescribed end-points and an intermediate keypoint. The players objective is respectively to maximize and minimize the path cost

$$\mathfrak{C}(\Xi, \Gamma) := \sup_{\xi \in \Xi} \inf_{\gamma \in \Gamma} \mathfrak{C}(\xi, \gamma), \quad \text{where } \mathfrak{C}(\xi, \gamma) := \int_0^{T(\gamma)} \mathcal{C}_\xi(\gamma(t), \gamma'(t), \gamma''(t)) dt, \quad (1)$$

where the path γ is parametrized at unit euclidean speed, and the final time $T(\gamma)$ is free. From a game theoretic point of view, this is a non-cooperative zero-sum game, where player Ξ has no information and player Γ has full information over the opponent's strategy.

The game (1) typically models a surveillance problem [15], and $\exp(-\mathfrak{C}(\Xi, \Gamma))$ is the probability for player Γ to visit the intermediate keypoint without being detected by player Ξ (by convention, players are identified with their sets of strategies). For instance player Γ could be a robber trying to steal an object from a museum, and player Ξ would be responsible for the placement of video surveillance cameras. In another instance, player Ξ is responsible for the installation of radar [1] or sonar detection systems [15], and would like to prevent vehicles sent by player Γ from spying on some objectives without being detected.

The dependence of the cost \mathcal{C}_ξ w.r.t. the path tangent $\gamma'(t)$ models the variation of a measure of how detectable the target is (radar cross section, directivity index, etc.) w.r.t. the relative positions and orientations of the target and sensor. The dependence of \mathcal{C}_ξ on the path curvature $\gamma''(t)$ models the airplane maneuverability constraints, such as the need to slow down in tight turns [6], or even a hard bound on the path curvature [5].

Strode [15] have shown the interplay of motion planning and game theory in a similar setting, on a multistatic sonar network use case, but using isotropic graph-based path planning. The same year, Barbaresco [2] used fast-marching for computing threatening paths toward a single radar, but without taking into account curvature constraints and without considering a game setting.

The main contributions of this paper are as follows:

1. *Anisotropy and curvature penalization:* Strategy optimization for player Γ is an optimal motion planning problem, with a known cost function. This is addressed by numerically solving a generalized eikonal PDE posed on a two or three dimensional domain, and which is strongly anisotropic in the presence of a curvature penalty and a detection measurement that depends on orientation. A Fast-Marching algorithm,

relying on recent adaptive stencils constructions, based on tools from lattice geometry, is used for that purpose [6, 10, 11]. In contrast, the classical fast marching method [13] used in [4] is limited to cost functions $\mathcal{C}_\xi(\gamma(t))$ independent of the path orientation $\gamma'(t)$ and curvature $\gamma''(t)$.

2. *Gradient computation for sensors placement:* Strategy optimization for player Ξ is typically a non-convex problem, to which various strategies can be applied, yet gradient information w.r.t. the variable $\xi \in \Xi$ is usually of help. For that purpose, we implement efficient differentiation algorithms, forward and reverse, for estimating the gradient of the value function of player Ξ

$$\nabla_\xi \mathfrak{C}(\xi, \Gamma), \quad \text{where } \mathfrak{C}(\xi, \Gamma) := \inf_{\gamma \in \Gamma} \mathfrak{C}(\xi, \gamma). \quad (2)$$

Reverse mode differentiation reduced the computation cost of $\nabla_\xi \mathfrak{C}(\xi, \Gamma)$ from $\mathcal{O}(N^2)$, as used in [4], to $\mathcal{O}(N \ln N)$, where N denotes the number of ization points of the domain. As a result, we can reproduce examples from [4] with computation times reduced by several orders of magnitude, and address complex three dimensional problems.

Due to space constraints, this paper is focused on problem modeling, ization and numerical experiments, rather than on mathematical aspects of wellposedness and convergence analysis. Free and open source codes for reproducing (part of) the presented numerical experiments are available on the first author's webpage¹.

2 Mathematical background of trajectory optimization

We describe in this section the PDE formalism, based on generalized eikonal equations, used to compute the value function $\min_{\gamma \in \Gamma} \mathfrak{C}(\xi, \gamma)$ of the second player, where ξ is known and fixed. Their discretization is discussed in §3. We distinguish two cases, depending on whether the path cost function $\mathcal{C}_\xi(x, \dot{x}, \ddot{x})$ appearing in (1) depends on the last entry \ddot{x} , i.e. on path curvature.

¹github.com/Mirebeau/HamiltonianFastMarching

2.1 Curvature independent cost

Let $\Omega \subseteq \mathbb{E} := \mathbb{R}^2$ be a bounded domain, and let the source set Υ and target set Θ be subsets of Ω . For each $x \in \Omega$, let Γ_x denote the set of all paths $\gamma \in C^1([0, T], \Omega)$, where $T = T(\gamma)$ is free, such that

$$\gamma(0) \in \Upsilon, \quad \gamma(T) = x, \quad \text{and } \forall t \in [0, T], \|\gamma'(t)\| = 1.$$

The problem description states that the first player needs to go from Υ to Θ and back, hence its set of strategies is $\Gamma = \bigcup_{x \in \Theta} \Gamma_x \times \Gamma_x$. Therefore

$$\mathfrak{C}(\xi, \Gamma) = 2 \inf_{x \in \Theta} u_\xi(x), \quad \text{where } u_\xi(x) := \inf_{\gamma \in \Gamma_x} \mathfrak{C}(\xi, \gamma). \quad (3)$$

Define the 1-homogenous metric $\mathcal{F}_\xi : \Omega \times \mathbb{E} \rightarrow [0, \infty]$, the Lagrangian \mathcal{L}_ξ and the Hamiltonian \mathcal{H}_ξ by

$$\mathcal{F}_\xi(x, \dot{x}) := \|\dot{x}\| \mathcal{C}_\xi(x, \dot{x}/\|\dot{x}\|), \quad \mathcal{L}_\xi := \frac{1}{2} \mathcal{F}_\xi^2, \quad \mathcal{H}_\xi(x, \hat{x}) := \sup_{v \in \mathbb{E}} \langle \hat{x}, v \rangle - \mathcal{L}_\xi(x, v).$$

Under mild assumptions [3], the function $u_\xi : \Omega \rightarrow \mathbb{R}$ is the unique viscosity solution to a generalized eikonal equation

$$\forall x \in \Omega \setminus \Upsilon, \mathcal{H}_\xi(x, \nabla_x u_\xi(x)) = 1/2, \quad \forall x \in \Upsilon, u_\xi(x) = 0,$$

with outflow boundary conditions on $\partial\Omega$. The discretization of this PDE is discussed in §3. We limit in practice our attention to Isotropic costs $\mathcal{C}_\xi(x)$, and Riemannian costs $\mathcal{C}_\xi(x, \dot{x}) = \sqrt{\langle \dot{x}, M_\xi(x) \dot{x} \rangle}$ where $M_\xi(x)$ is symmetric positive definite, for which efficient numerical strategies have been developed [8, 10].

2.2 Curvature dependent cost

Let $\Omega \subseteq \mathbb{R}^2 \times \mathbb{S}^1$ be a bounded domain, within the three dimensional space of all positions and orientations. As before, let $\Upsilon, \Theta \subseteq \Omega$. For all $x \in \Omega$ let Γ_x^\pm be the collection of all $\gamma \in C^2([0, T], \Omega)$, such that $\eta := (\gamma, \pm\gamma')$ satisfies

$$\eta(0) \in \Upsilon, \quad \eta(T) = x, \quad \text{and } \forall t \in [0, T], \|\gamma'(t)\| = 1.$$

Here and below, the symbol “ \pm ” must be successively replaced with “ $+$ ” and then “ $-$ ”. Since the first player needs to go from Υ to Θ and back, its set of strategies is $\Gamma = \bigcup_{x \in \Theta} \Gamma_x^+ \times \Gamma_x^-$. Therefore

$$\mathfrak{C}(\xi, \Gamma) = \inf_{x \in \Theta} u_\xi^+(x) + u_\xi^-(x), \quad \text{where } u_\xi^\pm(x) := \inf_{\gamma \in \Gamma_x^\pm} \mathfrak{C}^\pm(\xi, \gamma). \quad (4)$$

We denoted \mathfrak{C}^\pm the path cost defined in terms of the local cost $\mathcal{C}_\xi(p, \pm\dot{p}, \ddot{p})$. Consider the 1-homogeneous metric $\mathcal{F}_\xi^\pm : T\Omega \rightarrow [0, \infty]$, defined on the tangent bundle to $\Omega \subseteq \mathbb{R}^2 \times \mathbb{S}^1$ by

$$\mathcal{F}_\xi^\pm((p, n), (\dot{p}, \dot{n})) := \begin{cases} +\infty & \text{if } \dot{p} \neq \|\dot{p}\|n, \\ \|\dot{p}\|\mathcal{C}_\xi(p, \pm n, \dot{n}/\|\dot{p}\|) & \text{else,} \end{cases}$$

where $p \in \mathbb{R}^2$, $n \in \mathbb{S}^1$ is a unit vector, and the tangent vector satisfies $\dot{p} \in \mathbb{R}^2$, $\dot{n} \perp n$.

Introducing the Lagrangian $\mathcal{L}_\xi^\pm = \frac{1}{2}(\mathcal{F}_\xi^\pm)^2$ on $T\Omega$, and its Legendre-Fenchel dual the Hamiltonian \mathcal{H}_ξ^\pm , one can again under mild assumptions characterize u_ξ^\pm as the unique viscosity solution to the generalized eikonal PDE $\mathcal{H}_\xi^\pm(x, \nabla u_\xi^\pm(x)) = 1/2$ with appropriate boundary conditions [3]. In practice, we choose cost functions of the form $\mathcal{C}_\xi(p, \dot{p}, \ddot{p}) = \mathcal{C}_\xi^\circ(p, \dot{p})\mathcal{C}_*(|\ddot{p}|)$, where \mathcal{C}_* is the Reeds-Shepp car or Dubins car [5] curvature penalty, namely

$$\mathcal{C}_{\text{RS}}(\kappa) := \sqrt{1 + \rho^2 \kappa^2}, \quad \mathcal{C}_{\text{D}}(\kappa) := \begin{cases} 1 & \text{if } |\rho\kappa| \leq 1, \\ +\infty & \text{otherwise,} \end{cases}$$

where $\rho > 0$ is a parameter which has the dimension of a curvature radius. The Dubins car can only follows path which curvature radius is $\leq \rho$, whereas the Reeds-Shape car, in the sense of [6] and without reverse gear), can rotate into place if needed. The Hamiltonian then has the explicit expression $\mathcal{H}((p, n), (\hat{p}, \hat{n})) = \frac{1}{2}\mathcal{C}_\xi^0(p, n)^{-2}\mathcal{H}_*(n, (\hat{p}, \hat{n}))$ where

$$\mathcal{H}_{\text{RS}} = \frac{1}{2}(\langle \hat{p}, n \rangle_+^2 + \|\hat{n}/\rho\|^2) \quad \mathcal{H}_{\text{D}} = \frac{1}{2} \max\{0, \langle \hat{p}, n \rangle + \|\hat{n}/\rho\|\}^2.$$

3 Discretization of generalized eikonal equations

We construct a discrete domain X by intersecting the computational domain with an orthogonal grid of scale $h > 0$

$$X = \Omega \cap (h\mathbb{Z})^d,$$

where $d = 2$ for the curvature independent models, $d = 3$ for the other models which are posed on $\mathbb{R}^2 \times (\mathbb{R}/2\pi\mathbb{Z})$ —using the angular parametrization $\mathbb{S}^1 \cong \mathbb{R}/2\pi\mathbb{Z}$ (in the latter periodic case, and $2\pi/h$ must be an integer). We design weights $c_\xi(x, y)$, $x, y \in X$ such that for any tangent vector \dot{x} at x one has

$$\mathcal{H}_\xi(x, \dot{x}) \approx h^{-2} \sum_{y \in X} c_\xi^2(x, y) \langle x - y, \dot{x} \rangle_+^2, \quad (5)$$




Figure 1: Examples of adaptive stencils used by the Fast-Marching solver

where $a_+ := \max\{0, a\}$. (Expression (5), is typical although some models require a slight generalization.) The weights $c_\xi(x, y)$ are non-zero for only few $(x, y) \in X$ at distance $\|x - y\| = \mathcal{O}(h)$, see Figure 1. Their construction exploits the additive structure of the discretization grid X and relies on techniques from lattice geometry [14], see [6, 10, 11] for details. The generalized eikonal PDE $\mathcal{H}_\xi(x, \nabla_x u_\xi(x)) = 1/2$ is discretized as

$$\sum_{y \in X} c_\xi^2(x, y) (U_\xi(x) - U_\xi(y))_+^2 = h^2/2 \quad (6)$$

with adequate boundary conditions. The solution $U_\varepsilon : X \rightarrow \mathbb{R}$ to this system of equations is computed in a single pass with $\mathcal{O}(N \ln N)$ complexity using the Fast-Marching algorithm [13].

To be able to use the gradient to solve the problem (1), we need to differentiate the cost $\mathfrak{C}(\xi, \Gamma)$ w.r.t. the first player strategy $\xi \in \Xi$. In view of (3) and (4), this only requires the sensitivity of the discrete solution values $U_\varepsilon(x_*)$ at the few points $x_* \in X \cap \Theta$, w.r.t to variations in the weights $c_\varepsilon(x, y)$, $x, y \in X$. For that purpose we differentiate (6) w.r.t. ξ and obtain

$$\sum_{y \in X} \omega_\xi(x, y) (dU_\xi(x) - dU_\xi(y) + (U_\xi(x) - U_\xi(y)) d \ln c_\xi(x, y)) = 0,$$

where $\omega_\xi(x, y) := c_\xi^2(x, y) (U_\xi(x) - U_\xi(y))_+$. This equation ties $dU_\varepsilon(x)$ to those $dU_\xi(y)$, $y \in X$, such that $\omega_\xi(x, y) > 0$, hence satisfying $U_\xi(x) < U_\xi(y)$ by construction. In the spirit of automatic differentiation by reverse accumulation [16], these relations are back-propagated from any Θ keypoint x_* to the Υ seeds, at a modest cost $\mathcal{O}(N)$.

4 Numerical results

The chosen physical domain R is the rectangle $[0, 2] \times [0, 1]$ minus some obstacles, as illustrated on Figure 2. Source point is $(0.2, 0.5)$ and target




Figure 2: No intervention from the first player. Shortest path from the blue point (left) to the red keypoint (right) and back.

keypoint $(1.8, 0.5)$. The computational domain is thus $\Omega = R$ for curvature independent models and $\Omega = R \times \mathbb{S}^1$ for curvature dependent models, which is discretized on a 180×89 or $180 \times 89 \times 60$ grid.

No intervention from the first player. The cost function is

$$\mathcal{C}_\xi(p, \dot{p}, \ddot{p}) = \mathcal{C}_*(|\ddot{p}|), \quad (7)$$

where $\mathcal{C}_*(\kappa)$ is respectively 1 , $\sqrt{1 + \rho^2 \kappa^2}$ and $1 + \infty \times \chi_{\rho \kappa > 1}$, with $\rho := 0.3$. The differences between the three models are apparent: the curvature independent model uses the same path forward and back; the Reeds-Shepp car spreads some curvature along the way but still makes an angle at the target point; the Dubins car maintains the radius of curvature below the bound ρ , and its trajectory is a succession of straight and circular segments.

Fresh paint based detection. In this toy model, see Figure 3, the first player spreads some fresh paint over the domain, and the second player is regarded as detected if he comes back covered in it from his visit to the keypoint. The cost function is

$$\mathcal{C}_\xi(p, \dot{p}, \ddot{p}) = \xi(p) \mathcal{C}_*(|\ddot{p}|) \quad (8)$$

where $\xi : R \rightarrow \mathbb{R}_+$ is the fresh paint density, decided by the first player, and $\mathcal{C}_*(\kappa)$ is as above. For wellposedness, we impose a bounds on the density, namely $0.1 \leq \xi \leq 1$, and subtract the paint supply cost $\int_R \xi(p) dp$ to (1). The main interest of this specific game, also considered in [4], is that $\mathfrak{C}(\xi, \Gamma)$ is concave w.r.t. $\xi \in \Xi$. The optimal strategy for player Ξ is: in the curvature independent case to make some “fences” of paint between close obstacles, and in the curvature penalized models to deposit paint at the edges of obstacles, as well as along specific circular arcs for the Dubins model. At the optimum for player Ξ , a continuum of paths have equal cost for player Γ , covering a large portion of the domain with a uniform density as illustrated on Figure 3 (bottom), see [4] for details.

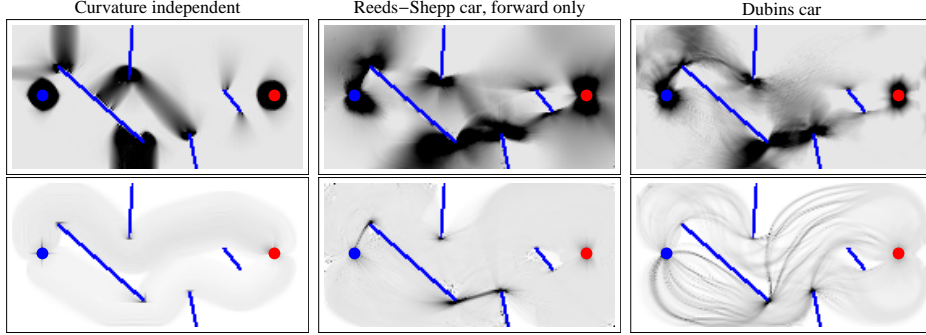


Figure 3: Fresh paint based detection. Top: Optimal distribution of paint, to mark a path from the blue point (left) to the red keypoint (right) and back. Bottom: Geodesic density at the equilibrium.

Visual detection. The first player places some cameras, which efficiency at detecting the second player decreases with distance and is blocked by obstacles, see Figure 4. The cost function is

$$\mathcal{C}_\xi(p, \dot{p}, \ddot{p}) = \mathcal{C}_*(\kappa) \sum_{\substack{q \in \xi \\ [p, q] \subseteq R}} \frac{1}{\|q - p\|^2}, \quad (9)$$

where $\xi \in \Xi$ is a subset of R with prescribed cardinality, two in our experiments. The green arrows on Fig 4 originate from the current (non optimal) camera position, and point in the direction of greatest growth $\nabla \mathfrak{C}(\xi, \Gamma)$ for the first player objective function.

Radar based detection. The first player places some radars on the domain $R = [0, 2] \times [0, 1]$, here devoid of obstacles, and the second player has to fly by undetected. The cost function is

$$\mathcal{C}_\xi(p, \dot{p}, \ddot{p}) = \mathcal{C}_*(|\ddot{p}|) \sqrt{\sum_{q \in \xi} \frac{\langle \dot{p}, n_{pq} \rangle^2 + \delta^2 \langle \dot{p}, n_{pq}^\perp \rangle^2}{\|p - q\|^4}} \quad (10)$$

where $n_{pq} := (q - p)/\|q - p\|$, and ξ is a three element subset of $[0.4, 1.6]$. The parameter δ is set to 1 for an isotropic radar cross section (RCS), or to 0.2 for an anisotropic RCS. In the latter where a plane showing its side to radar is five times less likely to be detected than a plane showing its nose or back, at the same position. Green arrows on Figure 5 point from the original position to the (locally) optimized position for player Ξ . At the equilibrium, several paths are optimal for player Γ , shown in red on Fig 5.

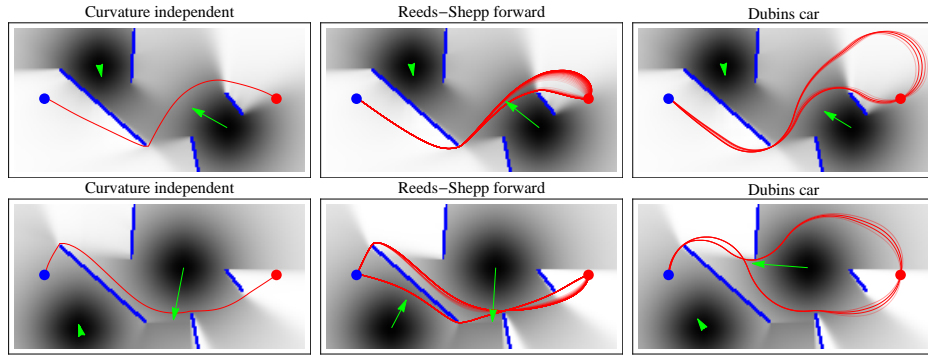


Figure 4: Visual detection. Field of view of the cameras (black gradients), optimal furtive paths (red lines), local direction of improvement of the camera position (green arrows).

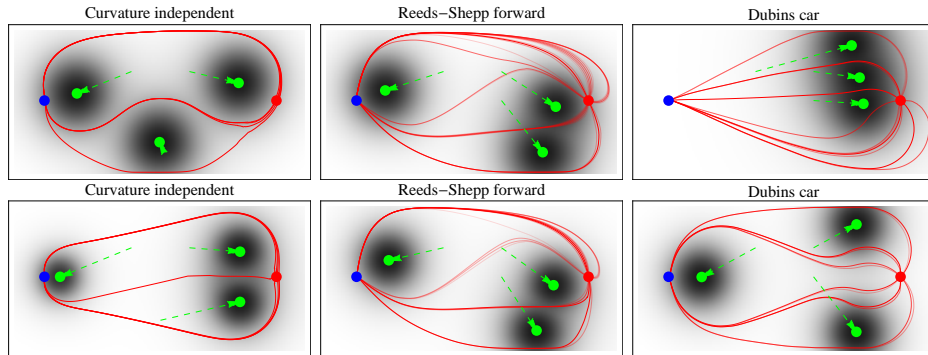


Figure 5: Radar based detection. Optimal radar placement with an isotropic (top) or anisotropic (bottom) radar cross section.

Computational cost On a standard Laptop computer (2.7Ghz, 16GB ram), optimizing the second player objective, by solving a generalized eikonal equation, takes ≈ 1 s in the curvature dependent case, and ≈ 60 times less in the curvature independent case thanks to the absence of angular discretization of the domain. Optimizing the first player objective takes ≈ 100 $L - BFGS$ iterations, each one taking at most 8s. For the stability of the minimization procedure, the problems considered were slightly regularized by the use of soft-minimum functions and by “blurring” the target keypoint over the 3×3 box of adjacent pixels.

5 Conclusion

We have modeled a motion planning problem that take into account navigation constraints and minimizing an anisotropic probability of detection, while computing the gradient of the value function related to the sensors location. This model is thus useful for surveillance applications modeled as a two-player zero-sum game involving a target that tries to avoid detection.

References

- [1] F. Barbaresco and B. Monnier, *Minimal geodesics bundles by active contours: Radar application for computation of most threatening trajectories areas & corridors*, 10th European Signal Processing Conference, Tampere, 2000, pp. 1–4.
- [2] F. Barbaresco, *Computation of most threatening radar trajectories areas and corridors based on fast-marching & Level Sets*, IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA), Paris, 2011, pp. 51–58.
- [3] M. Bardi and I. Capuzzo-Dolcetta, *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*, Birkhauser, 1997.
- [4] F. Benmansour, G. Carlier, G. Peyr, and F. Santambrogio, *Derivatives with respect to metrics and applications: subgradient marching algorithm*, Numerische Mathematik, vol. 116, no. 3, pp. 357381, May 2010.
- [5] L. E. Dubins, *On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents*, Amer. J. Math., vol. 79, pp. 497516, 1957.

- [6] R. Duits, S.P.L. Meesters, J.-M. Mirebeau, J. M. Portegies, *Optimal Paths for Variants of the 2D and 3D Reeds-Shepp Car with Applications in Image Analysis*, Preprint available on arXiv
- [7] J. Fehrenbach and J.-M. Mirebeau, *Sparse Non-negative Stencils for Anisotropic Diffusion*, Journal of Mathematical Imaging and Vision, pp. 125, 2013.
- [8] J.-M. Mirebeau, *Anisotropic Fast-Marching on cartesian grids using Lattice Basis Reduction*, SIAM J. Numer. Anal., vol. 52, no. 4, pp. 15731599, Jan. 2014.
- [9] J.-M. Mirebeau, *Minimal stencils for discretizations of anisotropic PDEs preserving causality or the maximum principle*, SIAM J. Numer. Anal., vol. 54, no. 3, pp. 15821611, 2016.
- [10] J.-M. Mirebeau, *Anisotropic fast-marching on cartesian grids using Voronois first reduction of quadratic forms*, in preparation
- [11] J.-M. Mirebeau, *Fast Marching methods for Curvature Penalized Shortest Paths*, in preparation
- [12] D. Mumford, *Elastica and Computer Vision*, no. 31, New York, NY: Springer New York, pp. 491506, 1994.
- [13] E. Rouy and A. Tourin, *A Viscosity Solutions Approach to Shape-From-Shading*, SIAM J. Numer. Anal., vol. 29, no. 3, pp. 867884, Jul. 1992.
- [14] A. Schürmann, *Computational geometry of positive definite quadratic forms*, University Lecture Series, 2009.
- [15] C. Strode, *Optimising multistatic sensor locations using path planning and game theory*, IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA), Paris, 2011, pp. 9–16.
- [16] A. Griewank, A. Walther, *Evaluating derivatives: principles and techniques of algorithmic differentiation*. Society for Industrial and Applied Mathematics, 2008.